

611-TD-591-001

EOSDIS Core System Project

M&O Procedures: Section 6—Network Administration

Interim Update

July 2001

Raytheon Company
Upper Marlboro, Maryland

Preface

This document is an interim update to the Mission Operations Procedures Manual for the ECS Project, document number 611-CD-600-001. This document has not been submitted to NASA for approval, and should be considered unofficial.

This document has been updated to reflect changes in the User and Production Networks and the resulting changes in DAAC LAN topology and terminology as well as changes resulting from the removal of DCE mechanisms and HP machines from the baseline.

Any questions should be addressed to:

Data Management Office
The ECS Project Office
Raytheon Company
1616 McCormick Drive
Upper Marlboro, Maryland 20774-5301

This page intentionally left blank.

6. Network Administration

This section covers the procedures necessary for the management operations that monitor and control the system network capabilities.

Detailed procedures for tasks performed by the Network Administrator are provided in the sections that follow. The procedures assume that the administrator is authorized and has proper access privileges to perform the tasks (i.e., root).

6.1 Network documentation

ECS Network Administration requires access to restricted documents not available via the cmdm.east.hitc.com (PETE Server) URL. These documents include:

- Network Overview Diagram 921-TDx-001
(x = DAAC designation: G/Goddard; L/Langley; N/NSIDC; E/EDC)
- Hardware Network Diagram 921-TDx-002
- Host IP Assignments 921-TDx-003
- Network IP Assignments 921-TDx-004
- Dual-Homed Host Static Routes 921-TDx-005
- Ingest Host Static Routes 921-TDx-006

These documents describe and depict the network layout and inter/intra-connections necessary to understand the ECS. Contact Landover Configuration Management for copies for individual sites.

6.2 Network Monitoring

Monitoring is conducted primarily via command-line interaction and site-developed scripts. An overview of network information is provided as a basis for evaluating network status.

6.3 DAAC LAN Topology Overview

The Distributed Active Archive Center (DAAC) Local Area Network (LAN) consists of a User Fiber Distributed Data Interface (FDDI) Network, a Production/Ingest Ethernet Network, and a High Performance Parallel Interface (HiPPI) Network. Figures 6.3-1 through 6.3-4 illustrate overviews of the LAN topology for the EROS Data Center (EDC) DAAC, the National Snow and Ice Data Center (NSIDC) DAAC, the Langley Research Center (LaRC) DAAC, and the Goddard Space Flight Center (GSFC) DAAC. As the figures show, there are variations in the topology at the different sites. Note: The NSIDC DAAC does not have Ingest or HiPPI networks.

The creation of separate User and Processing networks allows processing flows to be unaffected by user pull demands, and the introduction of the high-speed HiPPI Network provides adequate bandwidth to the Processing and Data Server subsystems to transfer high volumes of data. Each of the networks is discussed in detail below.

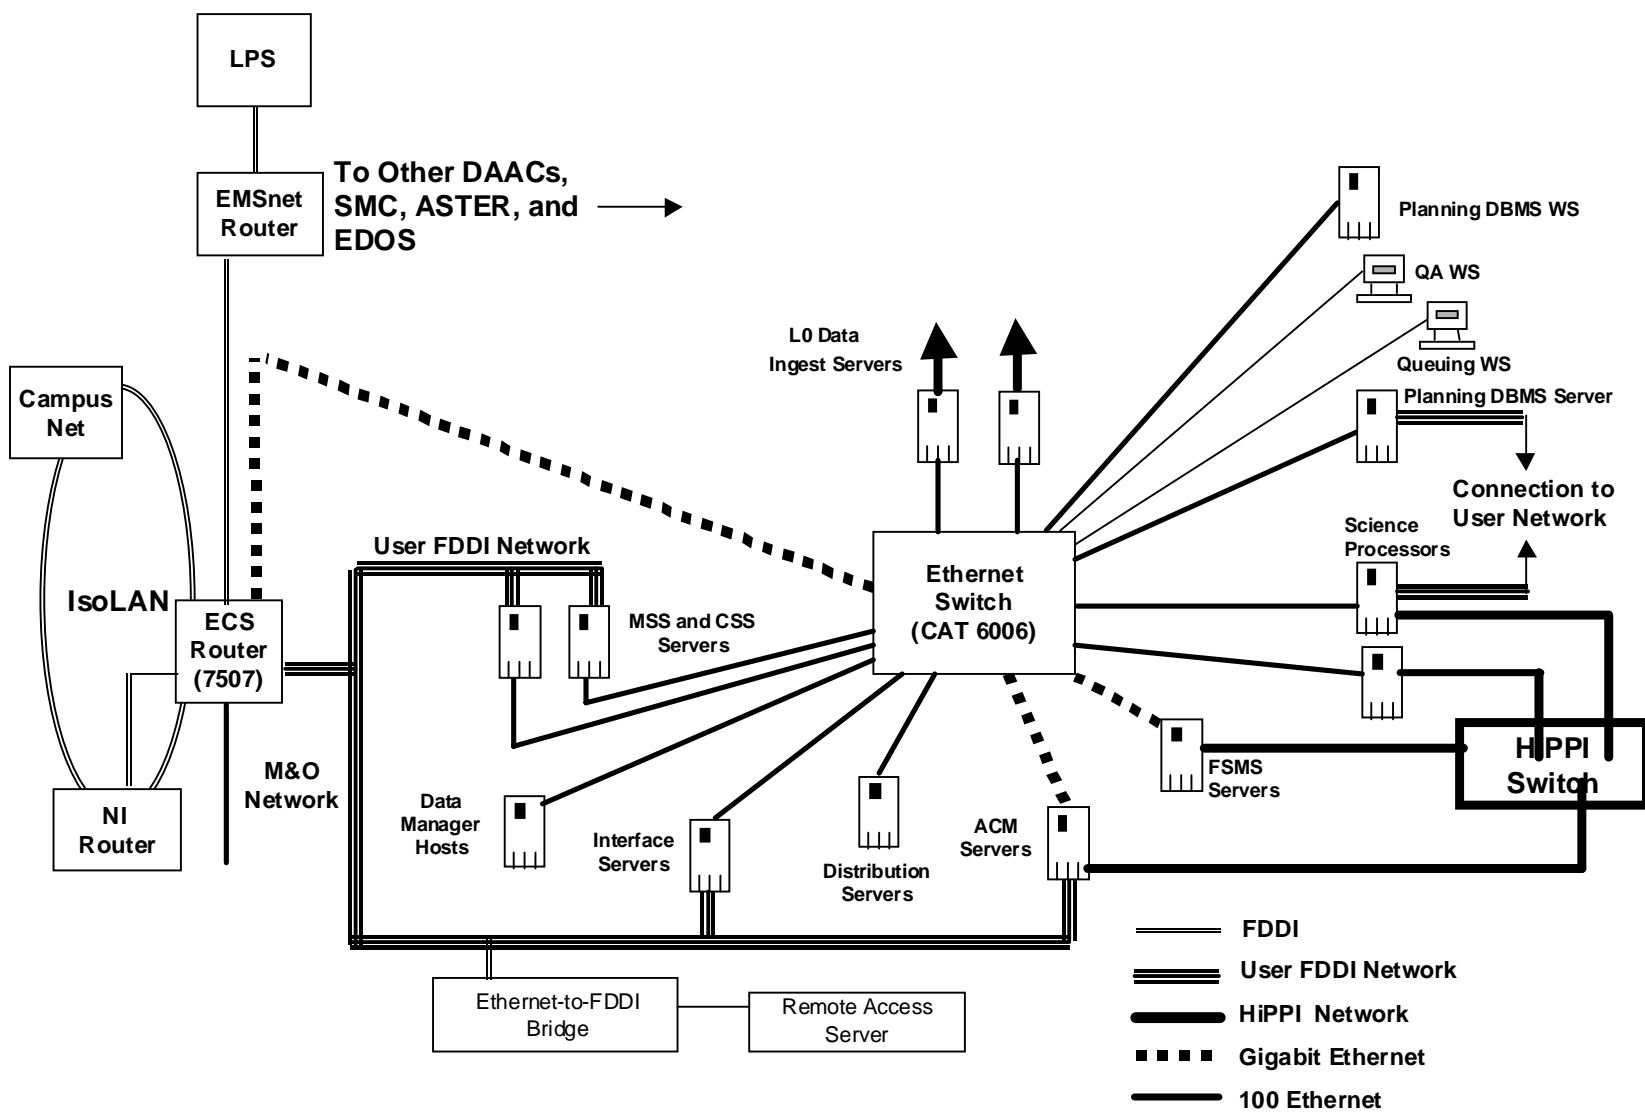


Figure 6.3-1. EDC DAAC LAN Topology

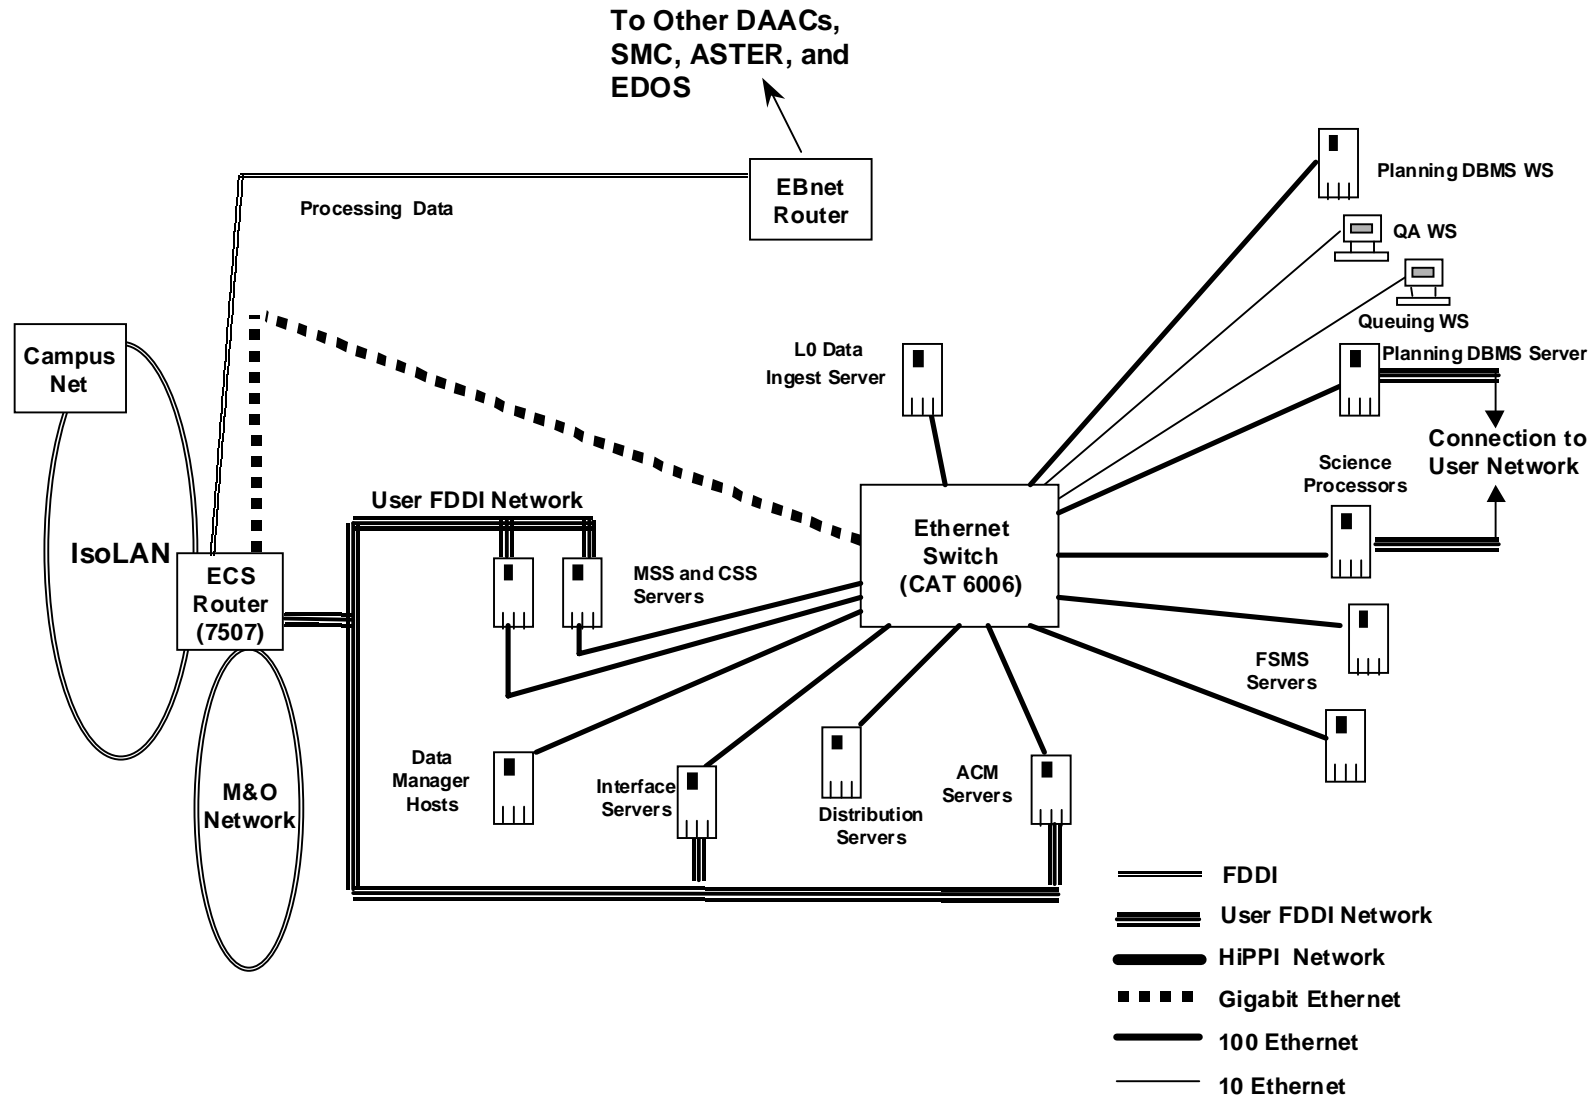


Figure 6.3-2. NSIDC DAAC LAN Topology

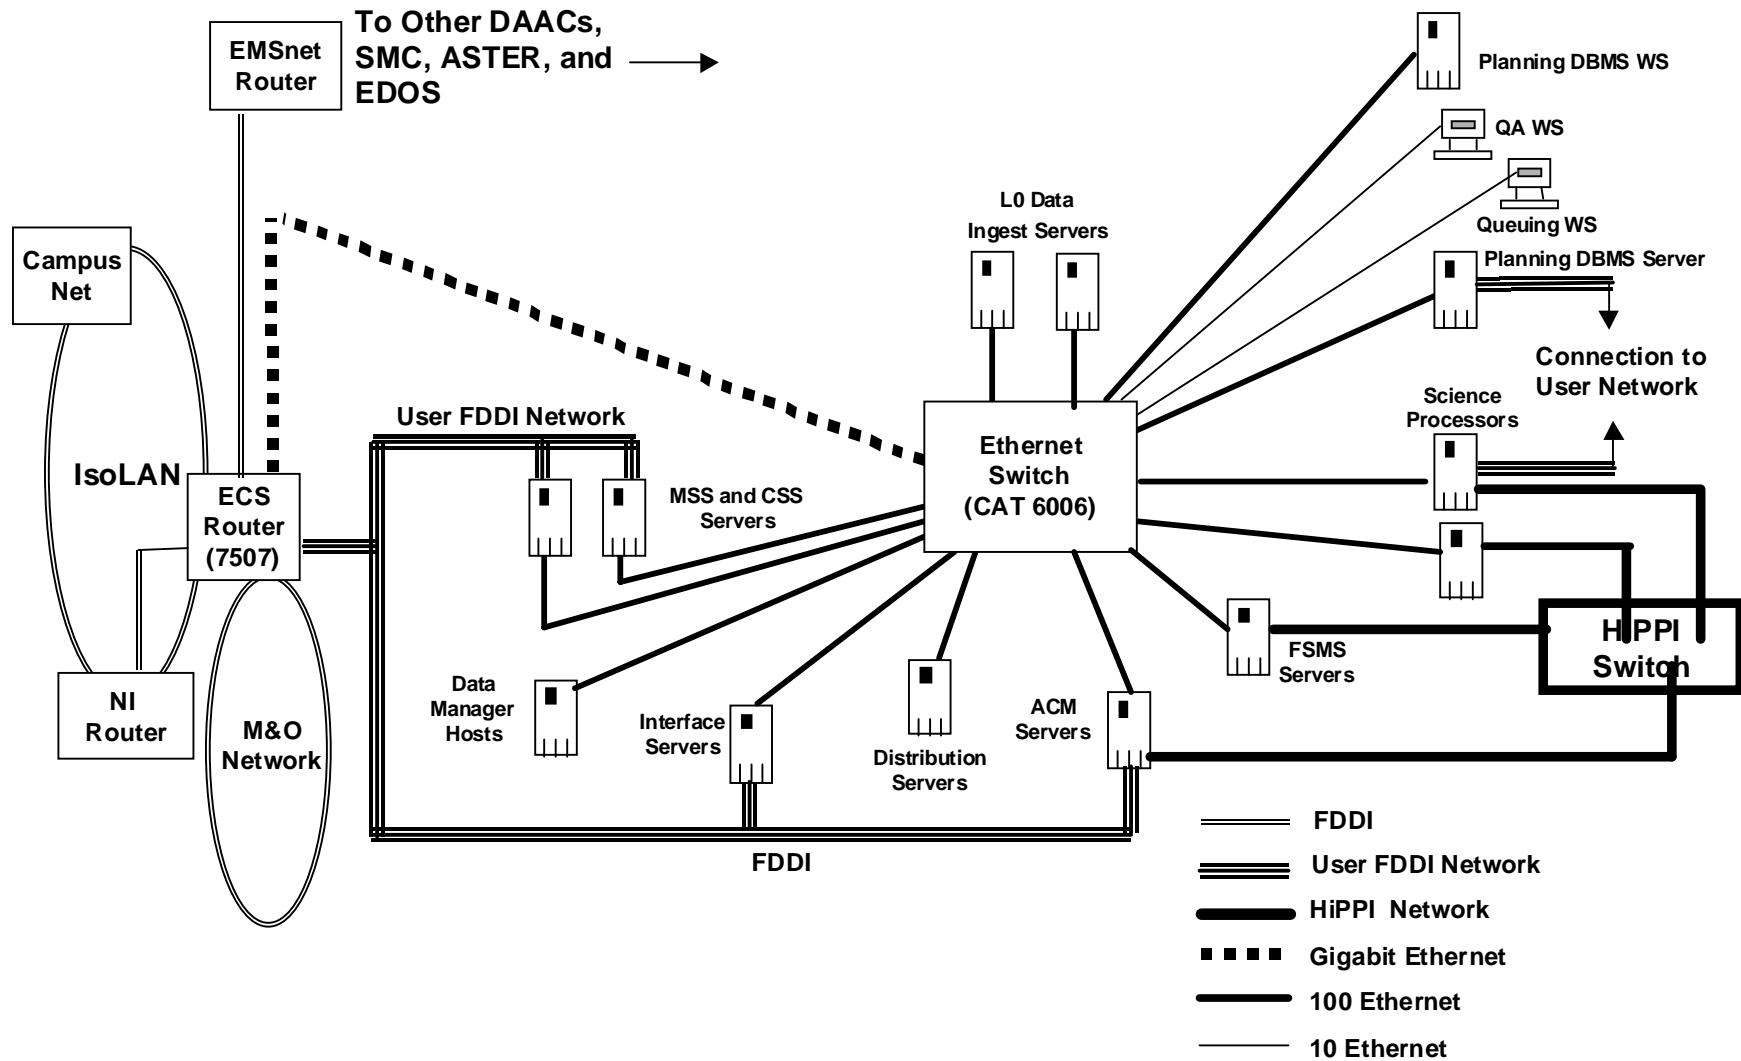


Figure 6.3-3. LaRC DAAC LAN Topology

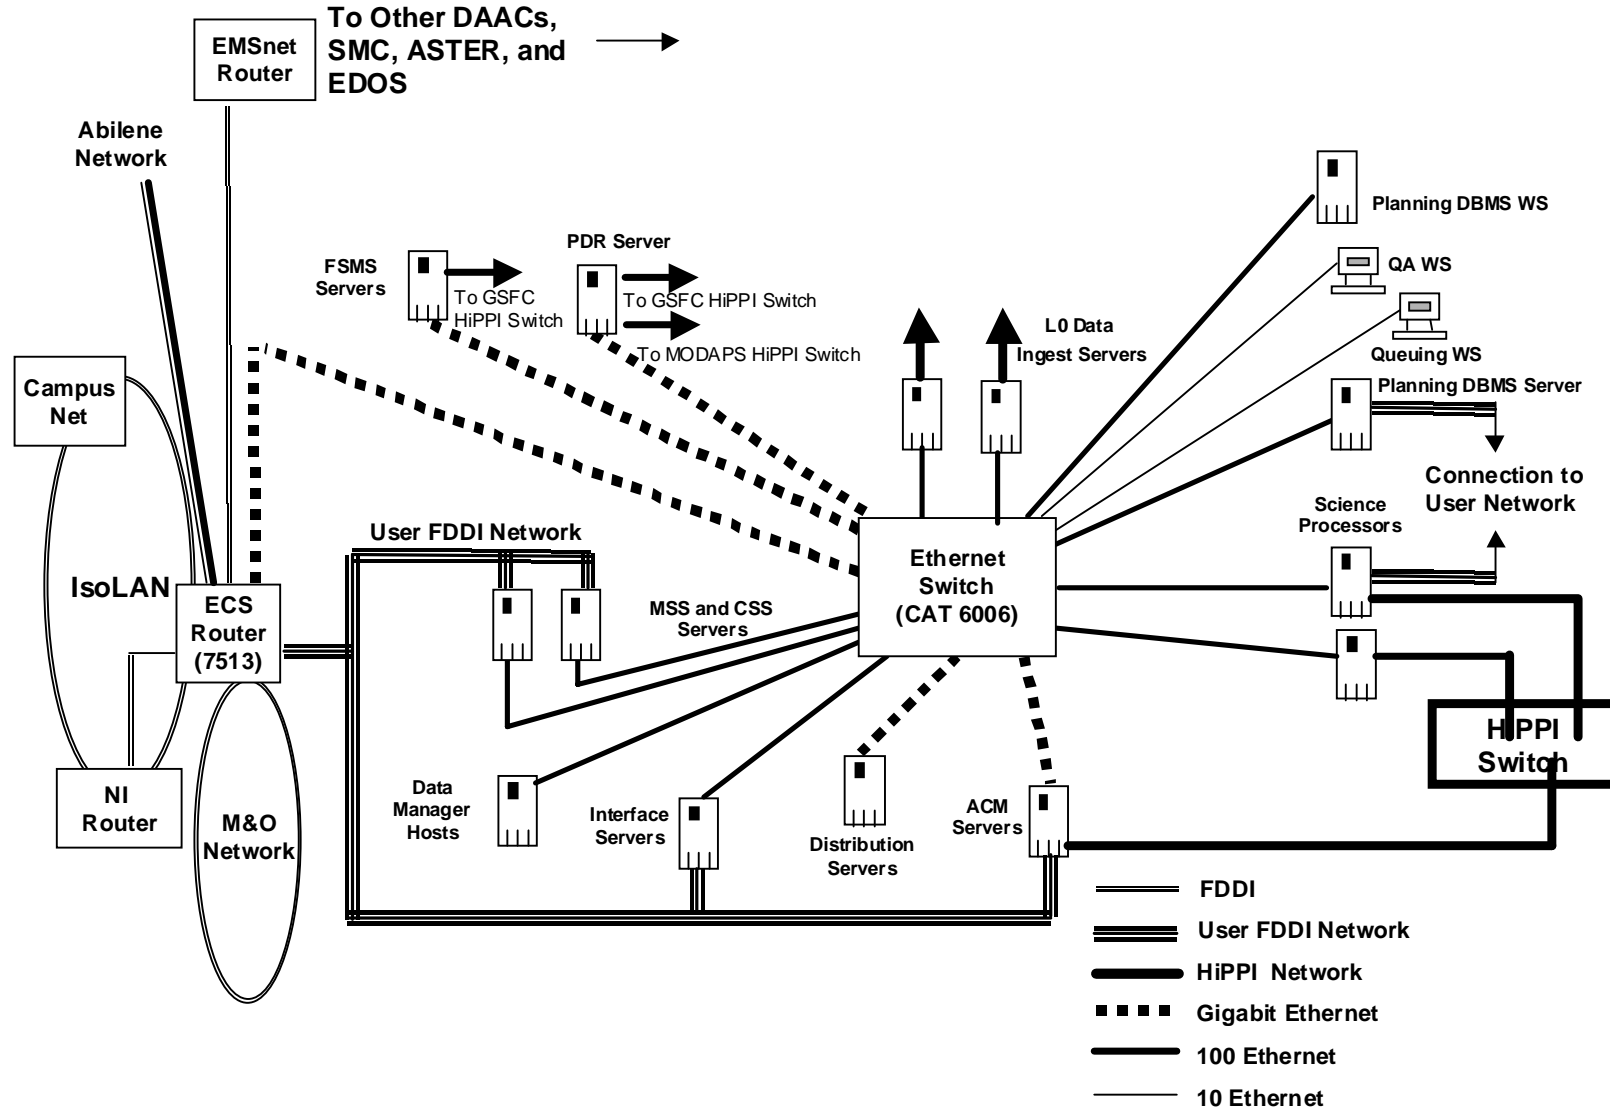


Figure 6.3-4. GSFC DAAC LAN Topology

6.3.1 The Production Network

The Production Network consists of a Catalyst 6006 Ethernet switch supporting the DAAC subsystems. EOS Mission Support network (EMSn) [formerly EOSDIS Backbone Network (Ebnet)] connections to external production systems such as EDOS and other ECS DAACs are made by means of the DAAC's ECS router. A connection in the ECS router provides access to the EMSn router to handle DAAC-DAAC flows.

6.3.2 The User Network

The User Network is an FDDI-based LAN connecting users (via NI, local campuses, general Internet, etc.) to DAAC hosts responsible for providing user access. It has the main advantage of separating user and production flows. This allows DAAC processing data flows to be unaffected by user demand, so that even unanticipated user pull will not hinder the production network. Basically, the User Network provides access to Data Manager hosts and to a subset of DataServer hosts that interact directly with users. Users will not have access to any other hosts such as Ingest or Processing devices. CSS and MSS servers are also connected to the User Network. These servers are required for communications with outside networks for such things as name lookups and receipt of Internet mail as well as communication with and monitoring of DAAC interfaces with the user community (e.g., NI and the local campus).

The following subsystems have servers attached to the User Network:

- Planning Server
- Science Processor (1)
- MSS Servers
- CSS Servers
- Interface Servers
- APC Servers

At EDC, there is also an Access Server attached to the User Network. This access server has two 28.8 Kbps modems attached so that remote users can access the EDC DAAC User Network via a dial-up circuit as well as via the Internet.

The User Network connects to the Campus Isolation LAN through an ECS router, which will provide the necessary routing and filtering controls. NI, the local Campus, and other Internet providers will also be connected to the Campus Isolation LAN. At some DAACs, there is an additional direct FDDI connection to the NI router because of the expected large user data flows.

6.3.4 Ethernet Topology

All hosts within a DAAC are connected to the Catalyst 6006 Ethernet switch. This switch is used to connect hosts at 10/100/100 Mb/s. The Catalyst 6006 Ethernet switch is also connected to the ECS router via a 1000Mb/s circuit.

6.3.5 FDDI Topology

User Network host interfaces have FDDI interface cards, which will be connected to a FDDI concentrator. This FDDI concentrator will be connected to the ECS Router.

6.3.6 The HiPPI Network

The HiPPI Network interconnects Data Server hosts and Science Processors in order to provide a high-speed network to handle the large data transfers between the two subsystems. The HiPPI network will be implemented via a central HiPPI switch with switched 800 Mbps interface ports connected directly to the high-powered processing and storage hosts. The HiPPI Network shifts the numerous transfers of large volumes of data onto a dedicated high-speed topology, freeing the Ethernet-based Production Network to handle control flows and DAAC-DAAC processing flows.

6.4 Network Hardware Components

6.4.1 LAN Components

The DAAC LANs consist of the following hardware components:

FDDI Concentrator. The FDDI Concentrator is a SynOptics 2914. It is a high-speed interface (100 Mbps). It consist of an FDDI interface for connecting to a FDDI ring and 12 FDDI ports for connecting to FDDI based hosts. It is only used on the User network.

Ethernet-to-FDDI Hub. The Ethernet-to-FDDI hub is a Cabletron MicroMAAC-22E. It consist of an FDDI interface for connecting to a FDDI ring and either 12 or 24 shared Ethernet ports. It is only used at EDC to connect the access server to the User network.

Access Server. The Access Server is a Cisco 2509. It consist of eight modem ports and an Ethernet port.

Maintenance and configuration of the access server is considered a non-trivial function. Such tasks are addressed in special technical training provided by the vendor and supplemental training provided by ECS. No further discussion of this device will be presented in this course.

ECS Router. The ECS Router is a Cisco 7507 or 7513. It is a high-speed interface (1000 Mbps). It consists of several FDDI and Ethernet interfaces. It interfaces to EMSnet, the local campus network, NI, M&O network, User network, and Production network. It provides IP address and port level filtering in support of the ECS security policy.

Maintenance and configuration of the ECS router is considered a non-trivial function. Such tasks are addressed in special technical training provided by the vendor and supplemental training provided by ECS. No further discussion of this device is presented in this course.

Ethernet Switch. The Ethernet switch is a Cisco catalyst 6006. It provides a large number of 10/100/100Mb/s interfaces. It interfaces to all Production hosts and to the ECS router. Maintenance and configuration of the Ethernet Switch is considered a non-trivial function. Such tasks are addressed in special technical training provided by the vendor and supplemental training provided by ECS. No further discussion of this device will be presented in this course.

HiPPI Switch. The HiPPI Switch is an Essential Communication EPS-16. It is a very high-speed interface (800 Mbps). It is used to connect some of the SGI hosts to a high speed switching fabric. Note: The NSIDC DAAC does not have a HiPPI Switch. Maintenance and

configuration of the HiPPI Switch is considered a non-trivial function. Such tasks are addressed in special technical training provided by the vendor and supplemental training provided by ECS. No further discussion of this device will be presented in this course.

6.5 ECS Domain Name Services (DNS) Structure

The parent DNS domain for ECS is `ecs.nasa.gov`. These DNS servers reside at the SMC, NSIDC, and EDC. In this domain are the SMC hosts, User hosts for all DAACs, and pointers to the DAACs' DNS servers.

The `ecs.nasa.gov` DNS servers are:

- `m0mss02.ecs.NASA.GOV` (Internet address = 198.118.212.37).
- `m0mss04.ecs.NASA.GOV` (Internet address = 198.118.212.41).
- `n0css02u.ecs.NASA.GOV` (Internet address = 198.118.206.84).
- `e0css02u.ecs.NASA.GOV` (Internet address = 198.118.203.104).

The DAACs' Production and HiPPI networks are a child domain of `ecs.nasa.gov`. They are:

- LaRC Production and HiPPI networks:
 - `l0ins02.larcb.ecs.nasa.gov` (Internet address = 198.118.219.74).
 - `l0css02.larcb.ecs.nasa.gov` (Internet address = 198.118.219.67).
- EDC Production and HiPPI networks:
 - `e0ins02.edcb.ecs.nasa.gov` (Internet address = 198.118.202.159).
 - `e0css02.edcb.ecs.nasa.gov` (Internet address = 198.118.202.132).
- NSIDC Production network:
 - `n0ins02.nsidcb.ecs.nasa.gov` (Internet address = 198.118.205.145).
 - `n0css02.nsidcb.ecs.nasa.gov` (Internet address = 198.118.205.123).
- GSFC Production and HiPPI networks:
 - `g0ins02.gsfc.ecs.nasa.gov` (Internet address = 198.118.210.69).
 - `g0css02.gsfc.ecs.nasa.gov` (Internet address = 198.118.210.63).

The DAACs' M&O networks are also a child domain of `ecs.nasa.gov`. They are:

- LaRC M&O network.
 - `larcmo.ecs.nasa.gov`.
- EDC M&O network.
 - `edcmo.ecs.nasa.gov`.
- NSIDC M&O network.
 - `nsidcmo.ecs.nasa.gov`.
- GSFC M&O network.
 - `gsfcmo.ecs.nasa.gov`.

6.6 Host Names

Because some hosts are attached to multiple networks (i.e., Production, User, and HiPPI), a letter is appended to the production host name to distinguish which interface (and IP address) a user is accessing.

As an example, a GSFC DAAC host named g0acg01.gsfc.nasa.gov is a host attached to the Production network. If that host was also attached to the HiPPI network, it would also be known as g0acg01h.gsfc.nasa.gov. If this host were also attached to the User network, it would also be known as g0acg01u.gsfc.nasa.gov.

6.7 Network Security

6.7.1 ECS Network Connectivity

The ECS network was designed to minimize unauthorized user access. Ingest network access at a DAAC is limited to its Level 0 data provider(s), the SMC, and hosts attached to the DAAC's Production and M&O networks. No local campus, Internet or other DAAC access is provided. Access to a DAAC's Production network is limited to the SMC, the DAAC's M&O network, and other DAACs. No local campus, Internet, or Level 0 data provider(s) access is provided. The ECS Network Connectivity matrix, shown in

, summarizes which networks can access hosts on another network.

6.7.2 Troubleshooting - Verifying Connectivity

One of the key reasons for failure of data access and transfer is an error or problem in system connectivity. This can be caused by a myriad of glitches such as incorrect/outdated lookup tables, incorrectly assigned IP addresses, missing default route and more. Besides checking individual host/server operation with various tools such as Tivoli and Whazzup??, you can use several command line entries to verify point-to-point communication between components.

There are three initial steps to help verify system connectivity. They include ensuring connectivity is authorized, determining if the Domain Name Service (DNS) is resolving host name and IP addresses correctly, and actively testing the connectivity by using the ping function.

Authorized connectivity can be determined by checking the ECS Network Connectivity matrix, Table 6-5.

Table 6-5. ECS Network Connectivity

| Network | GSFC Prod/ Ingest | GSFC User | LaRC Prod/ Ingest | LaRC User | EDC Prod/ Ingest | EDC User | NSIDC Prod/ Ingest | NSIDC User | VATC Prod/ Ingest | VATC User | PVC Prod/ Ingest | PVC User | SMC | NI/ Internet |
|--------------------------------|-------------------------|--------------|-------------------------|--------------|------------------------|-------------|--------------------------|---------------|-------------------------|--------------|------------------------|-------------|-----|-----------------|
| GSFC Production/ Ingest | X | | X | | X | | X | | X | | X | | X | |
| GSFC User | | | | | | | | | | | | | X | X |
| LaRC Production/ Ingest | X | | X | | X | | X | | X | | X | | X | |
| LaRC User | | | | | | | | | | | | | X | X |
| EDC Production/ Ingest | X | | X | | X | | X | | X | | X | | X | |
| EDC User | | | | | | | | | | | | | X | X |
| NSIDC Production/ Ingest | X | | X | | X | | X | | X | | X | | X | |
| NSIDC User | | | | | | | | | | | | | X | X |
| VATC Production/ Ingest | X | | X | | X | | X | | X | | X | | X | |
| VATC User | | | | | | | | | | | | | X | X |
| PVC Production/ Ingest | X | | X | | X | | X | | X | | X | | X | |
| PVC User | | | | | | | | | | | | | X | X |
| SMC | X | X | X | X | X | X | X | X | X | X | X | X | X | X |
| NI/Internet | | X | | X | | X | | X | | X | X | X | X | |

Checking local host access to another local host over the network

- 1 On workstation **x0xxx##**, at the UNIX prompt in a terminal window, check the Domain Name Service entries (DNS) for the source host by typing **nslookup <local_host>**.
 - The screen display will be similar to the following:

```
g0spg01{mblument}[204]->nslookup g0spg01
Server: g0css02.gsfc.nasa.gov
Address: 198.118.210.63
Name: g0spg01.gsfc.nasa.gov
Address: 198.118.210.16
```
- 2 Check the DNS entries for the remote host by typing **nslookup <other_host>**.
 - The screen display will be similar to the following:

```
g0spg01{mblument}[201]->nslookup g0css02
Server: g0css02.gsfc.nasa.gov
Address: 198.118.210.63
Name: g0css02.gsfc.nasa.gov
Address: 198.118.210.63
```

- 3 Determine the host's network interface using **ifconfig <interface>** where **<interface>** parameter can be found by executing **netstat -i**

- The **netstat -i** command will provide the following information:
g0spg01{mblument}[201]->netstat -i

| Name | Mtu | Network | Address | Ipkts | Ierrs | Opkts | Oerrs | Coll |
|------|-------|----------------|----------------|----------|-------|----------|-------|------|
| ipg0 | 4352 | 198.118.210 | g0spg01.gsfc. | 9182666 | 1 | 8103032 | 0 | 0 |
| hip0 | 65280 | 192.168.1 | g0spg01h.gsfc. | 5554524 | 0 | 6776651 | 0 | 0 |
| xpi0 | 4352 | 198.118.212.64 | g0spg01u.ecs. | 37850320 | 0 | 14109683 | 3 | 0 |
| xpi1 | 0 | none | none | 0 | 0 | 0 | 0 | 0 |
| et0* | 1500 | none | none | 0 | 0 | 0 | 0 | 0 |
| lo0 | 8304 | loopback | localhost | 314800 | 0 | 314800 | 0 | 0 |
- Using **ipg0** from the **ifconfig <interface>** data as the interface parameter, **ifconfig ipg0**, will result in the following display:
g0spg01{mblument}[203]->ifconfig ipg0
ipg0: flags=863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST>
inet 198.118.210.16 netmask 0xfffff00 broadcast 198.118.210.255

- 4 Ping the two hosts to verify their inter-connectivity.

- Ping the local host (g0spg01).
g0spg01{mblument}[232]->ping g0spg01
PING g0spg01.gsfc.ecs.nasa.gov (198.118.210.16): 56 data bytes
64 bytes from 198.118.210.16: icmp_seq=0 ttl=255 time=0 ms
64 bytes from 198.118.210.16: icmp_seq=1 ttl=255 time=0 ms
64 bytes from 198.118.210.16: icmp_seq=2 ttl=255 time=0 ms
64 bytes from 198.118.210.16: icmp_seq=3 ttl=255 time=0 ms
64 bytes from 198.118.210.16: icmp_seq=4 ttl=255 time=0 ms
----g0spg01.gsfc.ecs.nasa.gov PING Statistics----
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
g0spg01{mblument}[233]->
- Ping the remote host (g0css02).
g0spg01{mblument}[202]->ping g0css02
PING g0css02.gsfc.ecs.nasa.gov (198.118.210.63): 56 data bytes
64 bytes from 198.118.210.63: icmp_seq=0 ttl=255 time=2 ms
64 bytes from 198.118.210.63: icmp_seq=1 ttl=255 time=1 ms
64 bytes from 198.118.210.63: icmp_seq=2 ttl=255 time=1 ms
64 bytes from 198.118.210.63: icmp_seq=3 ttl=255 time=1 ms
64 bytes from 198.118.210.63: icmp_seq=4 ttl=255 time=1 ms
----g0css02.gsfc.ecs.nasa.gov PING Statistics----
5 packets transmitted, 5 packets received, 0% packet loss

round-trip min/avg/max = 1/1/2 ms

- 5 Check the health of the interface by executing **netstat -i**, looking for Ierrs and/or Oerrs that if present (1 or 2 errors are ok, 100 are not ok) indicate an interface problem; check the syslog for any startup or logged problems from the OS.

| g0spg01{mblument}[218]->netstat -i | | | | | | | | |
|------------------------------------|-------|----------------|----------------|----------|-------|----------|-------|------|
| Name | Mtu | Network | Address | Ipkts | Ierrs | Opkts | Oerrs | Coll |
| ipg0 | 4352 | 198.118.210 | g0spg01.gsfc. | 9197317 | 1 | 8113487 | 0 | 0 |
| hip0 | 65280 | 192.168.1 | g0spg01h.gsfc. | 5554541 | 0 | 6776668 | 0 | 0 |
| xpi0 | 4352 | 198.118.212.64 | g0spg01u.ecs. | 37851779 | 0 | 14109837 | 3 | 0 |
| xpi1 | 0 | none | none | 0 | 0 | 0 | 0 | 0 |
| et0* | 1500 | none | none | 0 | 0 | 0 | 0 | 0 |
| lo0 | 8304 | loopback | localhost | 325510 | 0 | 325510 | 0 | 0 |

- 6 Check the routing table for accuracy and completeness by executing **netstat -rn**.

- The resultant display will be similar to the following:

g0spg01{mblument}[226]->netstat -rn

Routing tables

Internet:

| Destination | Gateway | Netmask | Flags | Refs | Use | Interface |
|----------------|----------------|--------------|-------|------------|-------|-----------|
| default | 198.118.212.65 | | UGS | 1 14060556 | | xpi0 |
| 127.0.0.1 | 127.0.0.1 | | UH | 7 270097 | | lo0 |
| 192.168.1 | 192.168.1.1 | 0xffffffff00 | U | 0 | 0 | hip0 |
| 192.168.1.1 | 192.168.1.1 | | UGHS | 0 | 22 | hip0 |
| 192.168.1.2 | 192.168.1.2 | | UGHS | 1 3993577 | | hip0 |
| 192.168.1.3 | 192.168.1.3 | | UGHS | 0 | 17 | hip0 |
| 192.168.1.4 | 192.168.1.4 | | UGHS | 0 2397593 | | hip0 |
| 192.168.1.5 | 192.168.1.5 | | UGHS | 2 | 178 | hip0 |
| 192.168.1.6 | 192.168.1.6 | | UGHS | 0 | 24 | hip0 |
| 192.168.1.7 | 192.168.1.7 | | UGHS | 0 | 1403 | hip0 |
| 192.168.1.8 | 192.168.1.8 | | UGHS | 0 | 6 | hip0 |
| 192.168.1.9 | 192.168.1.9 | | UGHS | 0 | 0 | hip0 |
| 192.168.1.10 | 192.168.1.10 | | UGHS | 0 | 0 | hip0 |
| 198.118.198 | 198.118.210.1 | 0xffffffff00 | UGS | 0 | 16 | ipg0 |
| 198.118.198.12 | 198.118.210.2 | | UGHD | 0 | 31646 | ipg0 |
| 198.118.198.14 | 198.118.210.2 | | UGHD | 0 | 1032 | ipg0 |
| 198.118.198.17 | 198.118.210.2 | | UGHD | 0 | 0 | ipg0 |
| 198.118.198.25 | 198.118.210.2 | | UGHD | 0 | 194 | ipg0 |
| 198.118.198.26 | 198.118.210.2 | | UGHD | 0 | 36153 | ipg0 |
| 198.118.198.27 | 198.118.210.2 | | UGHD | 0 | 23425 | ipg0 |
| 198.118.198.28 | 198.118.210.2 | | UGHD | 4 | 11686 | ipg0 |
| 198.118.198.29 | 198.118.210.2 | | UGHD | 0 | 1682 | ipg0 |
| 198.118.198.30 | 198.118.210.2 | | UGHD | 3 | 14760 | ipg0 |

```

198.118.198.32 198.118.210.2      UGHD      2  917384 ipg0
198.118.198.42 198.118.210.2      UGHD      0   87381 ipg0
198.118.198.76 198.118.210.2      UGHD      3  568062 ipg0
198.118.198.100 198.118.210.2      UGHD      0   1223 ipg0
198.118.198.107 198.118.210.2      UGHD      0    299 ipg0
198.118.198.113 198.118.210.2      UGHD      0    893 ipg0
198.118.198.116 198.118.210.2      UGHD      0   9438 ipg0
198.118.202    198.118.210.1  0xffffffff UGS      0     0 ipg0
198.118.205    198.118.210.1  0xffffffff UGS      0     0 ipg0
198.118.208    198.118.210.1  0xffffffff UGS      0     0 ipg0
198.118.210    198.118.210.16  0xffffffff U      177 5624292 ipg0
198.118.210.16 127.0.0.1      UGHS      15  55462 lo0
198.118.211.32 198.118.210.1  0xffffffe0 UGS      0   6842 ipg0
198.118.212.32 198.118.210.1  0xffffffe0 UGS      0   1004 ipg0
198.118.212.40 198.118.210.2      UGHD      0   6205 ipg0
198.118.212.64 198.118.212.69  0xffffffe0 U      0   4485 xpi0
198.118.212.160 198.118.210.1  0xffffffe0 UGS      0     0 ipg0
198.118.219    198.118.210.1  0xffffffff UGS      0     0 ipg0
198.118.220    198.118.210.1  0xffffffff UGS      0     0 ipg0
198.118.232    198.118.210.1  0xffffffff UGS      0    143 ipg0
210.138.100    198.118.210.1  0xffffffff UGS      0     0 ipg0
224            198.118.210.16  0xf0000000 US      0     2 ipg0
g0spg01{mblument}[227]->

```

- Ping the default IP address to ensure connectivity to the default route (default: 198.118.212.65)

```
g0spg01{mblument}[228]->ping 198.118.212.65
```

```
PING 198.118.212.65 (198.118.212.65): 56 data bytes
```

```
64 bytes from 198.118.212.65: icmp_seq=0 ttl=255 time=1 ms
```

```
64 bytes from 198.118.212.65: icmp_seq=1 ttl=255 time=1 ms
```

```
64 bytes from 198.118.212.65: icmp_seq=2 ttl=255 time=1 ms
```

```
64 bytes from 198.118.212.65: icmp_seq=3 ttl=255 time=1 ms
```

```
64 bytes from 198.118.212.65: icmp_seq=4 ttl=255 time=1 ms
```

```
64 bytes from 198.118.212.65: icmp_seq=5 ttl=255 time=1 ms
```

```
64 bytes from 198.118.212.65: icmp_seq=6 ttl=255 time=1 ms
```

```
----198.118.212.65 PING Statistics----
```

```
7 packets transmitted, 7 packets received, 0% packet loss
```

```
round-trip min/avg/max = 1/1/1 ms
```


- 7 Check the other host using the same steps.
- 8 Check other hosts using the same infrastructure components as the two hosts with the problem.
- 9 If the host you are trying to communicate with is attached to the Ethernet Hub, make sure that the “Don’t Fragment” bit in the IP header is NOT set on the host that is FDDI attached. The Ethernet Hub does not support MTU discovery so it will not inform the host that the packet is too big. It silently discards the packet. By default, the Sun hosts are improperly configured. Check the file /etc/init.d/inetinit to ensure that the command to reset the “Don’t Fragment” bit is included: **ndd -set /dev/ip ip_path_mtu_discovery 0**

Checking local host access to another local host over the HiPPI Switch

- 1 On workstation **x0xxx##**, at the UNIX prompt in a terminal window, check the Domain Name Service entries (DNS) for the source host by typing **nslookup <local_host>**.
 - The screen display will be similar to the following:

```
g0drg07{mblument}[205]->nslookup g0drg07h
Server: g0css02.gsfc.nasa.gov
Address: 198.118.210.63
Name: g0drg07h.gsfc.nasa.gov
Address: 192.168.1.6
```
- 2 Check the DNS entries for the remote host by typing **nslookup <other_host>**.
 - The screen display will be similar to the following:

```
g0drg07{mblument}[203]->nslookup g0spg01h
Server: g0css02.gsfc.nasa.gov
Address: 198.118.210.63
Name: g0spg01h.gsfc.nasa.gov
Address: 192.168.1.1
```
- 3 Determine the host’s network interface using **ifconfig <interface>** where **<interface>** parameter can be found by executing **netstat -i**
 - The **netstat -i** command will provide the following information:

```
g0drg07{mblument}[216]->netstat -i
Name Mtu Network Address Ipks Ierrs Opks Oerrs Coll
ipg0 4352 198.118.210 g0drg07.gsfc. 114659 0 63268 0 0
hip0 65280 192.168.1 g0drg07h.gsfc. 85 0 115 24 0
et0* 1500 none none 0 0 0 0 0
lo0 8304 loopback localhost 10062 0 10062 0 0
g0drg07{mblument}[217]->
```
 - Using **hip0** from the **ifconfig <interface>** data as the interface parameter, **ifconfig hip0**, will result in the following display:

```
g0drg07{mblument}[217]->ifconfig hip0
hip0: flags=1061<UP,NOTRAILERS,RUNNING,CKSUM>
inet 192.168.1.6 netmask 0xfffff00
```

g0dr07{mblument}[218]->

4 Ping the two hosts and the hip0 interface IP address to verify their inter-connectivity.

- Ping the local host (g0dr07h).
g0dr07{mblument}[205]->ping 192.168.1.1
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=255 time=1 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=255 time=1 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=255 time=1 ms
64 bytes from 192.168.1.1: icmp_seq=3 ttl=255 time=1 ms
----192.168.1.1 PING Statistics----
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max = 1/1/1 ms
- Ping the remote host (g0spg01h).
g0dr07{mblument}[205]->ping 192.168.1.1
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=255 time=1 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=255 time=1 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=255 time=1 ms
64 bytes from 192.168.1.1: icmp_seq=3 ttl=255 time=1 ms
----192.168.1.1 PING Statistics----
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max = 1/1/1 ms
- Ping the HiPPI switch interface (g0dr07h.gsfc.nasa.gov).
g0dr07{mblument}[204]->ping 192.168.1.6
PING 192.168.1.6 (192.168.1.6): 56 data bytes
64 bytes from 192.168.1.6: icmp_seq=0 ttl=255 time=1 ms
64 bytes from 192.168.1.6: icmp_seq=1 ttl=255 time=1 ms
64 bytes from 192.168.1.6: icmp_seq=2 ttl=255 time=1 ms
64 bytes from 192.168.1.6: icmp_seq=3 ttl=255 time=1 ms
64 bytes from 192.168.1.6: icmp_seq=4 ttl=255 time=1 ms
64 bytes from 192.168.1.6: icmp_seq=5 ttl=255 time=1 ms
----192.168.1.6 PING Statistics----
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max = 1/1/1 ms
g0dr07{mblument}[205]->

5 Check the health of the interface by executing **netstat -i**, looking for Ierrs and/or Oerrs that, if present, (1 or 2 errors are ok, 100 are not ok) indicate an interface problem; check the syslog for any startup or logged problems from the OS.

```

g0drg07{mblument}[211]->netstat -i
Name Mtu Network Address IpKts Ierrs OpKts Oerrs Coll
ipg0 4352 198.118.210 g0drg07.gsfc. 128799 0 70050 0 0
hip0 65280 192.168.1 g0drg07h.gsfc. 103 0 133 24 0
et0* 1500 none none 0 0 0 0 0
lo0 8304 loopback localhost 11822 0 11822 0 0
g0drg07{mblument}[212]->

```

6 Check the routing table for accuracy and completeness by executing **netstat -rn**.

- The resultant display will be similar to the following:

```
g0drg07{mblument}[212]->netstat -rn
```

Routing tables

Internet:

| Destination | Gateway | Netmask | Flags | Refs | Use | Interface |
|-----------------|----------------|--------------|-------|------|-------|-----------|
| default | 198.118.210.1 | | UGS | 2 | 241 | ipg0 |
| 127.0.0.1 | 127.0.0.1 | | UH | 4 | 4659 | lo0 |
| 192.168.1 | 192.168.1.6 | 0xffffffff00 | U | 0 | 0 | hip0 |
| 192.168.1.1 | 192.168.1.1 | | UGHS | 0 | 24 | hip0 |
| 192.168.1.2 | 192.168.1.2 | | UGHS | 0 | 24 | hip0 |
| 192.168.1.3 | 192.168.1.3 | | UGHS | 0 | 33 | hip0 |
| 192.168.1.4 | 192.168.1.4 | | UGHS | 0 | 0 | hip0 |
| 192.168.1.5 | 192.168.1.5 | | UGHS | 0 | 0 | hip0 |
| 192.168.1.6 | 192.168.1.6 | | UGHS | 0 | 17 | hip0 |
| 192.168.1.7 | 192.168.1.7 | | UGHS | 0 | 52 | hip0 |
| 192.168.1.8 | 192.168.1.8 | | UGHS | 0 | 0 | hip0 |
| 192.168.1.9 | 192.168.1.9 | | UGHS | 0 | 0 | hip0 |
| 192.168.1.10 | 192.168.1.10 | | UGHS | 0 | 0 | hip0 |
| 198.118.198.25 | 198.118.210.2 | | UGHD | 0 | 612 | ipg0 |
| 198.118.198.109 | 198.118.210.2 | | UGHD | 0 | 7395 | ipg0 |
| 198.118.210 | 198.118.210.45 | 0xffffffff00 | U | 36 | 64343 | ipg0 |
| 198.118.210.45 | 127.0.0.1 | | UGHS | 21 | 7892 | lo0 |
| 198.118.212.40 | 198.118.210.2 | | UGHD | 0 | 865 | ipg0 |
| 224 | 198.118.210.45 | 0xf0000000 | US | 0 | 2 | ipg0 |

```
g0drg07{mblument}[213]->
```

7 Check the other HiPPI attached hosts using the same steps.

8 Check other hosts using the same infrastructure components as the two hosts with the problem.

Checking host communication across EMSn

- 1 Check the DNS by executing `nslookup < local_host>` and `nslookup <other host>`.
- 2 Check the host route table using `netstat -r`.
- 3 Run `tracert <target host ip address>` or similar tool to discover which router or local route table is in error or not having sufficient route information.
- 4 Check the Route Advertisement diagram, ECS Connectivity Matrix, and the Network Security Design to see that the filters are not blocking communications or provide no path between hosts. Details are in the configurations of the FDDI switch or ECS router. Also check host TCP Wrappers.

Checking host attached to User Network communicating with Hosts on the Campus Network(s) or to the Internet

- 1 Check the DNS by executing `nslookup < local_host>` and `nslookup <other host>`.
- 2 Check the host route table using `netstat -rn`. Make sure that the `route_add` script has executed correctly. (Note: The `route_add` script will NOT execute correctly if the User Network interface or Cisco ECS router are not working correctly).
- 3 Check that the default route is pointing to the Cisco ECS router User Network interface and not pointing to the Production Network interface.
- 4 Check the Route Advertisement diagram, ECS Connectivity Matrix, and the Network Security Design to see that the filters are not blocking communications or provide no path between hosts. Details are in the configurations of the Ethernet switch or ECS router. Also check host TCP Wrappers.

Checking host communicating with any other host over the HiPPI Switch

- 1 Check the DNS by executing `nslookup < local_host>` and `nslookup <other host>`.
- 2 Check on the extent of the problem by pinging itself over the HiPPI interface, `ping 192.168.x.x` where x.x can be discovered by **`netstat -i`**. Do the same for other hosts connected to the HiPPI Switch; if all fail, it's a HiPPI switch problem (check the other supplied troubleshooting tip handout); if all other hosts are ok, continue.
- 3 If the HiPPI cabling has been removed (for example, for SGI maintenance) check for bent pins at the SGI and re-seat cable; if problem not cleared, go to the other troubleshooting tip handout to begin testing with the `hiptest` testing.
- 4 Make sure the HiPPI interface is up by executing **`hipcntl startup`**.
- 5 Check on the HiPPI interface status by executing **`hipcntl hipp0 status`** and scan results for error conditions (see SGI ref #1 below, table 3-2).
- 6 Check that the HiPPI interface is "pingable" by executing **`ping <host_ip_address>`**; if no response, check the interface for being UP by executing **`ifconfig hip0`**; if interface is UP, there is probably a cable or switch interface problem.
- 7 Check the SGI HiPPI interface error counts by executing **`netstat -i`**.
- 8 Check for the static routes in the host route table by executing **`netstat -rn`**. It should look similar to following:

| | | | | | |
|-------------|-------------|----------------|----|-----|------|
| 192.168.2 | 192.168.2.5 | 0xffffffff00 U | 0 | 0 | hip0 |
| 192.168.2.1 | 192.168.2.1 | UGHS | 0 | 9 | hip0 |
| 192.168.2.2 | 192.168.2.2 | UGHS | 0 | 0 | hip0 |
| 192.168.2.3 | 192.168.2.3 | UGHS | 0 | 212 | hip0 |
| 192.168.2.4 | 192.168.2.4 | UGHS | 0 | 0 | hip0 |
| 192.168.2.5 | 192.168.2.5 | UGHS | 0 | 12 | hip0 |
| 192.168.2.6 | 192.168.2.6 | UGHS | 19 | 60 | hip0 |
| 192.168.2.7 | 192.168.2.7 | UGHS | 0 | 7 | hip0 |
| 192.168.2.8 | 192.168.2.8 | UGHS | 0 | 0 | hip0 |

Check the HiPPI interface to switch loop by executing **hiptest -I 0x0100000x** where x is the HiPPI switch's port number of the host's connection to it.

Check other host interfaces to see if any others are working properly.

Check the switch counters by logging (telnet) into the HiPPI switch from a Production network host, and look at counters by executing **show counters all all**. Also check the state of the interfaces by executing **show state all all** and look for interface(s) with a "true" power-up initialization error, a "false" source interconnect, or a "false" destination interconnect.

6.7.3 Specific Security Limitations

In addition to limiting network access as described above, access is further limited by port level filters installed in the ECS router. In addition to the port filters, a host's tcp wrappers will further limit network access.

Note: Any service that is not listed below is an allowable service.

The following services are NOT permitted into a DAAC's Production and User networks:

- remote login (tcp port 513).
- remote shell (tcp port 514).
- telnet to hosts (tcp port 23).
- NFS (udp and tcp ports 2049).
- Port Mapper [RPC] (udp and tcp ports 111).
- Access to udp and tcp ports 255-1023 on NIS servers.
- 7. X-11* (udp and tcp ports 6000-6003). X-11 11 is a special case. While by default X-11 is not allowed to X servers (X-terminals), a DAAC can decide to allow X-11 access between a selected set of hosts within the DAAC and an external entity such as a remote SSI&T host or a host at another DAAC. This access would be granted by modifying the appropriate router filter tables.

Also, ftp is NOT allowed to the MSS and CSS servers from the User network (tcp port 21).

Each DAAC has its own M&O Sustaining Engineering network. Hosts attached to this network are NOT permitted to use the following services when communicating with their Production and Ingest networks:

6.8.1 Script Locations

There is a separate route add script for each host type (Sun and SGI). The scripts are located in the following directories:

- Sun:
 - script S87route_add is in directory /etc/rc2.d.
- SGI:
 - script S87route_add is in directory /etc/init.d with a soft link to /etc/rc2.d/S87route_add.